Computer Vision

Kenneth.Dawson-Howe@scss.tcd.ie

School of Computer Science and Statistics,

Trinity College, University of Dublin. Ireland .

Module contents

Week	Beginning	Folder	Topic	Format	**Mini Tests (10%)**	Live Classes	Practical Component
		Introduction	What is Computer Vision	Live		Lecture Tuesday @2pm	
1	12/09/2022	Introduction	Module Organisation	Live		Lecture Wednesday @1pm	Install OpenCV & TIPS.
		Images	Camera Model & Digitisation	Pre-recorded			Experiment with TIPS
		Images	Colour Image models	Pre-recorded			**Assignment 1 (30%), Part 1**
2	19/09/2022	Images	Noise & Smoothing	Pre-recorded	Images before Tuesday @12noon	Q&A Session Tuesday @2pm	
		Binary	Thresholding	Pre-recorded		Tutorial Wednesday @1pm	
		Binary	Cleaning Binary Images	Pre-recorded	Binary before Tuesday @12noon		
		Regions	Connectivity & Connected Component Analy	Pre-recorded			**Assignment 1, Part 2**
3	26/09/2022	Regions	kMeans clustering	Pre-recorded	Regions before Tuesday @12noon	Q&A Session Tuesday @2pm	
		Geometric	Transformations	Pre-recorded		Tutorial Wednesday @1pm	
		Geometric	Interpolation	Pre-recorded	Geometric before Tuesday @12noon		
		Histograms	1D & 3D histograms	Pre-recorded			**Assignment 1, Part 3**
4	03/10/2022	Histograms	Equalisation, Comparison & Back Projection	Pre-recorded	Histograms before Tuesday @12noon	Q&A Session Tuesday @2pm	
		Video	Introduction	Pre-recorded		Tutorial Wednesday @1pm	
		Video	Background Models	Pre-recorded	Video before Tuesday @12noon		
		Edges	Edge detection	Pre-recorded			**Assignment 1, Part 4**
5	10/10/2022	Edges	Contour segmentation	Pre-recorded		Q&A Session Tuesday @2pm	
		Edges	Hough transform	Pre-recorded		Tutorial Wednesday @1pm	
		Edges	Least squared error & RANSAC	Pre-recorded	Edges before Tuesday @12noon		
		Learning & Evaluation	Learning & Ground Truth	Pre-recorded			**Assignment 1, Part 5**
6	17/10/2022	Learning & Evaluation	Performance	Pre-recorded	Learning & Evaluation before Tuesday @12noon	Q&A Session Tuesday @2pm	
		Recognition	SPR	Pre-recorded		Tutorial Wednesday @1pm	
		Recognition	SPR Features	Pre-recorded	Recognition before Tuesday @12noon		
		Regions II	Watershed segmentation	Pre-recorded			
		Regions II	Mean shift	Pre-recorded	Regions II before Tuesday @12noon	Q&A Session Tuesday @2pm	
8	31/10/2022	Recognition II	Template Matching	Pre-recorded		Tutorial Wednesday @1pm	
		Recognition II	Chamfer Matching	Pre-recorded			Submit Assignment 1
		Recognition II	SVM	Pre-recorded	Recognition II before Tuesday @12noon		- before 2-Nov @11:59pm
		Recognition III	Haar	Pre-recorded			**Assignment 2 (10%)**
9	07/11/2022	Recognition III	Haar	Pre-recorded		Q&A Session Tuesday @2pm	- Sample Exam Question
		Recognition III	PCA	Pre-recorded	Recognition III before Tuesday @12noon	Tutorial Wednesday @1pm	- released on 9-Nov @5:30pm
		Video II	Tracking - Mean Shift	Pre-recorded			
10	14/11/2022	Video II	Tracking - Mean Shift	Pre-recorded		Q&A Session Tuesday @2pm	Submit Assignment 3
		Video II	Tracking - Optical Flow	Pre-recorded	Video II before Tuesday @12noon	Tutorial Wednesday @1pm	- before 20-Nov @11:59pm
		Features	Corner detection - Moravec	Pre-recorded			
11	21/11/2022	Features	Corner detection - Harris	Pre-recorded		Q&A Session Tuesday @2pm	
		Features	SIFT Keypoints	Pre-recorded		Tutorial Wednesday @1pm	
		Features	SIFT Matching & Recognition	Pre-recorded	Features before Tuesday @12noon		
12	28/11/2022	Conclusions	How to address vision problems	Live		Lecture Tuesday @2pm	
		Conclusions	Sample exam paper tutorial	Live		Tutorial Wednesday @1pm	

Based on *A Practical Introduction to Computer Vision with OpenCV* by Kenneth Dawson-Howe © Wiley & Sons Inc. 2014

Introduction – Human Vision

- When it was discovered that an image is projected inside the eye…
  - Scientists thought vision was solved.
  - But there is no homunculus in our brain.
- So how do we understand the images from the eyes?
  - 20% of the brain does “visual only” processing.
  - 40%+ of the brain does “vision+” processing.
    - Motor, attention, navigation, …
- What does the brain do?
  - Visual cortex – basic image (pre)processing
  - Dedicated areas – face recognition, motion analysis, …
  - We know very little really.

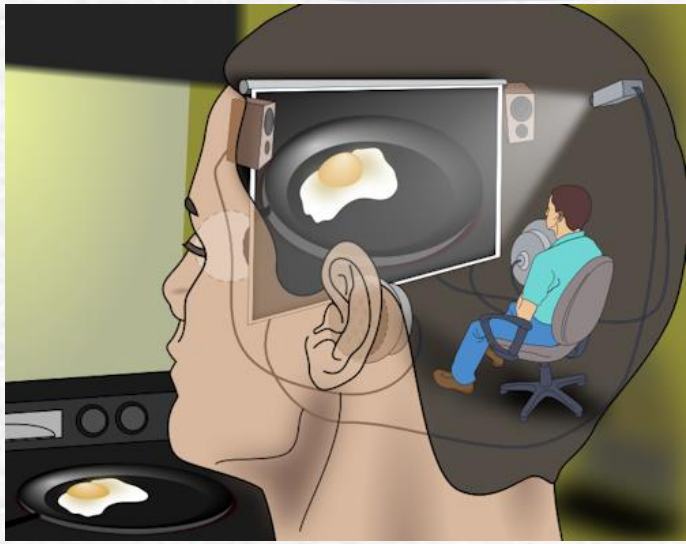

SOME RIGHTS RESERVED

Courtesy of Pbroks13 and Jennifer Garcia, see Wikipedia article on Cartesian theatre for more information

Introduction – Emulating human vision

- Machines can easily solve many tasks that we find difficult.
  - Particularly if we can describe them mathematically or in terms of an algorithm
- The solutions of many tasks that we find easy seem very difficult to describe formally.
  - Object recognition.
    - Rigid objects (variety)
    - Deformable objects
  - World modelling.

- Is machine vision just implementing human vision on a machine?
  - We don't understand human vision properly.
  - Machine vision is not restricted to two images from the visible spectrum.
    - UV, IR, Multiple cameras, Laser range finders, …
  - Machine learning can "tackle tasks that are too difficult to solve with fixed programs written and designed by human beings." From *Deep Learning* by Goodfellow

Based on *A Practical Introduction to Computer Vision with OpenCV* by Kenneth Dawson-Howe © Wiley & Sons Inc. 2014

Computer Vision is hard

"What we experience, apparently directly, is actually very different from what is recorded by our sense organs."

[Perception: From Sense to Object by J. Wilding]

Introduction

- Computer Vision is about understanding images. These images can be
  - Greyscale or Colour or Multi-spectral
  - Snapshots or video sequences
  - Taken with a static or moving camera
  - Taken of a stationary or dynamic scene
  - Taken with a calibrated or un-calibrated camera

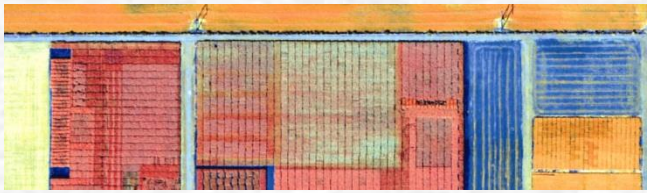

- What Computer Vision aims to do is to extract some useful information from these images for…
  - Inspection purposes
  - Analysis purposes
  - Control purposes

Based on *A Practical Introduction to Computer Vision with OpenCV* by Kenneth Dawson-Howe © Wiley & Sons Inc. 2014

Applications – Industrial Inspection

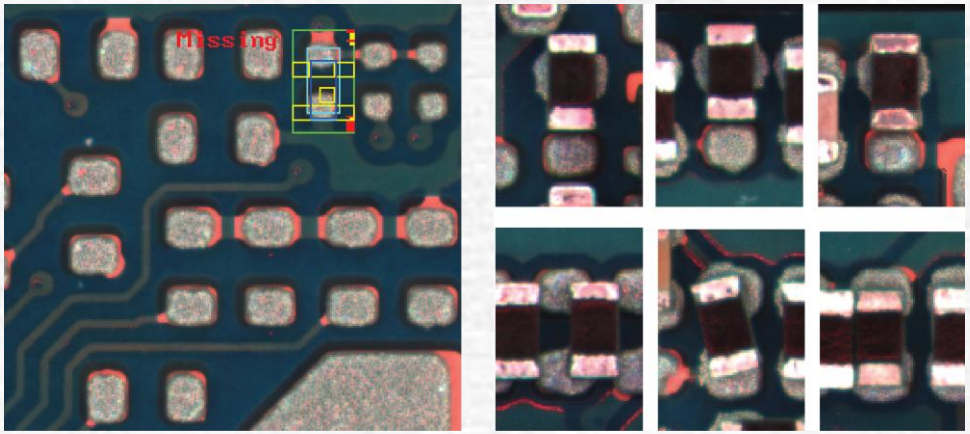

© Reproduced by permission of James Mahon

© Reproduced by permission of Omron Electronics LLC

Applications – Surveillance / Forensics

- In London there are 500,000+ cameras
- However, only 1 crime solved per 1,000 cameras (BBC, 2009)
- In July 2005 there was a major terrorist attack in London…
- Afterwards
  - 17,000 hours of video was studied
  - Over 1 year

© Image courtesy of Feelart / FreeDigitalPhotos.net

Applications – Biometrics

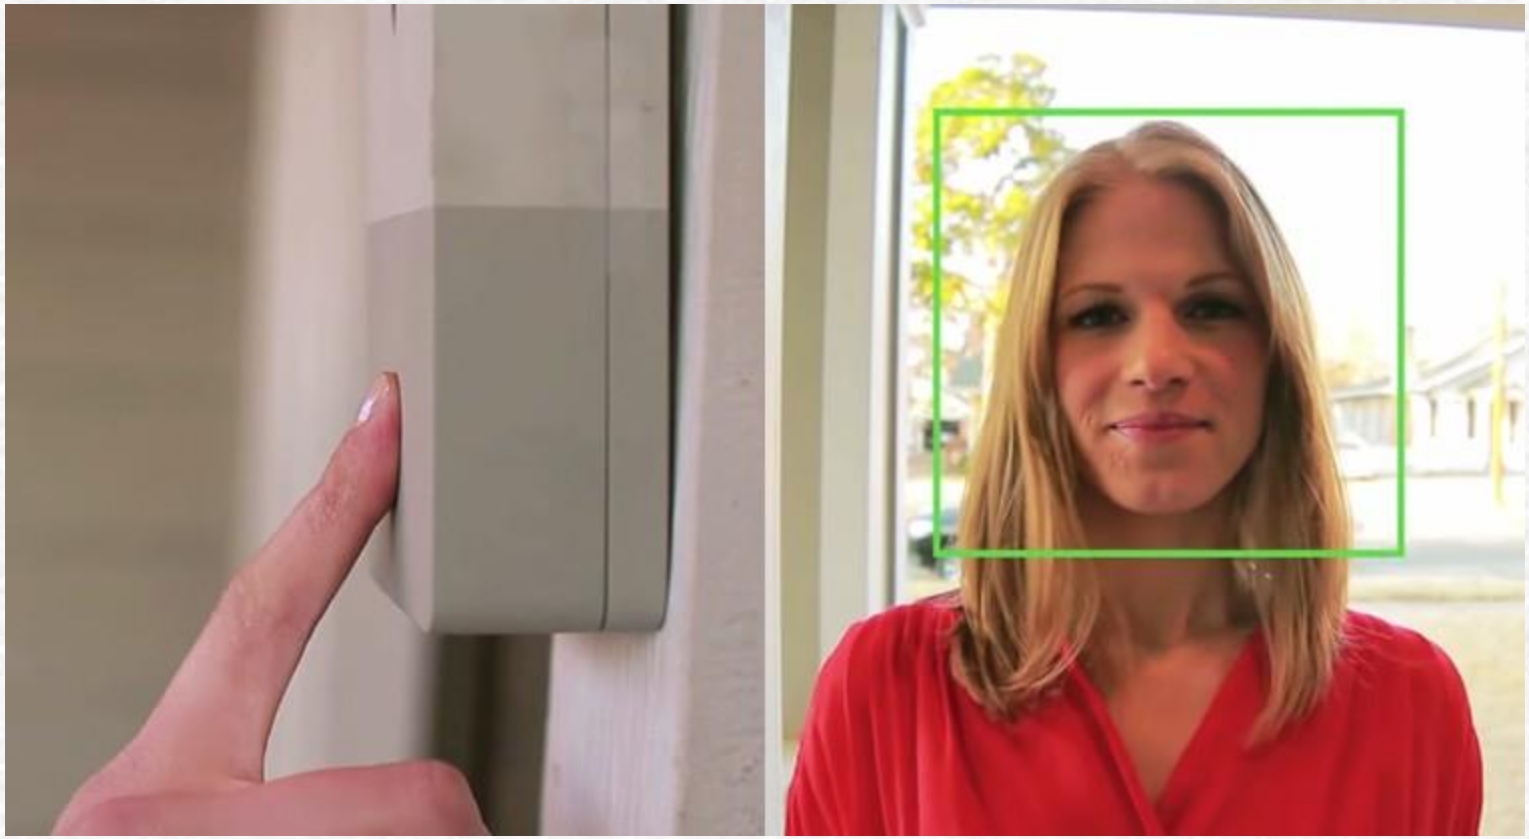

No need for a key. Just your face. Won't be fooled by photos.

Chui Smart Doorbell from http://iothought.com/chui-smart-doorbell/

Applications – Landmine detection

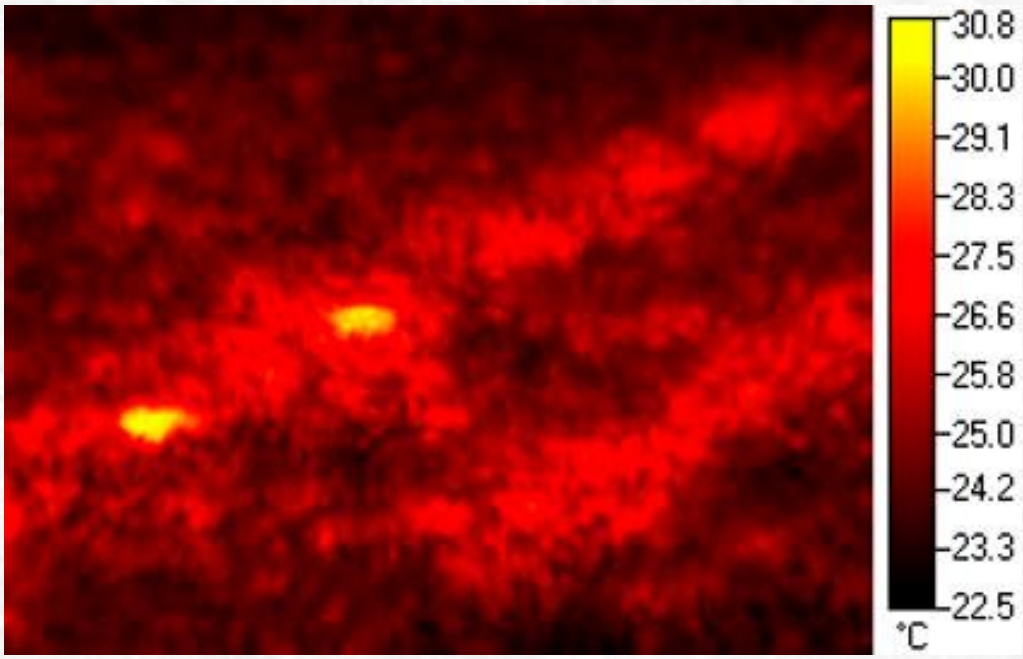

© Reproduced by permission of Zouheir Fawaz

Applications – Medical Imaging

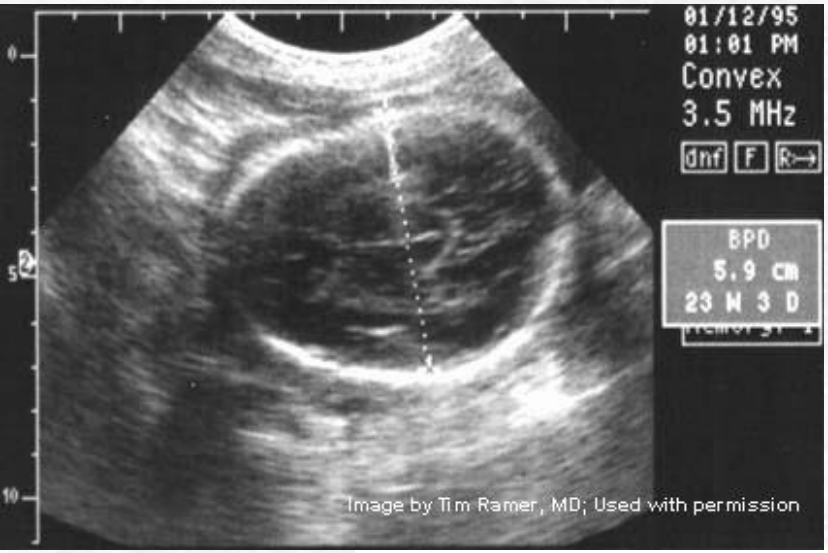

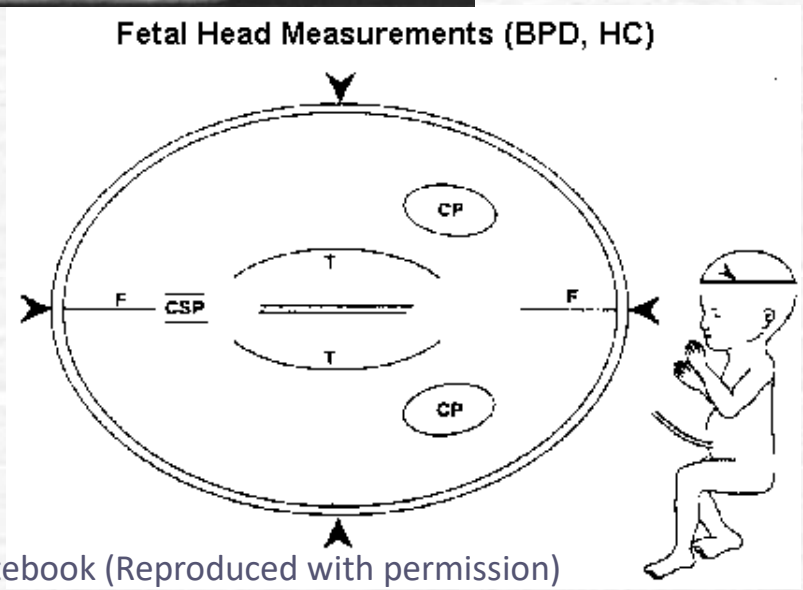

© Family Practice Notebook (Reproduced with permission)

© Etan J. Tal, released under Creative Commons 3.0

Vision is confusing

"Every image is the image of a thing merely for him who knows how to read it, and who is enabled by the aid of the image to form an idea of the thing."

[Handbook of Physiological Optics by H. Helmholtz]

Based on *A Practical Introduction to Computer Vision with OpenCV* by Kenneth Dawson-Howe © Wiley & Sons Inc. 2014

Interfacing with the world

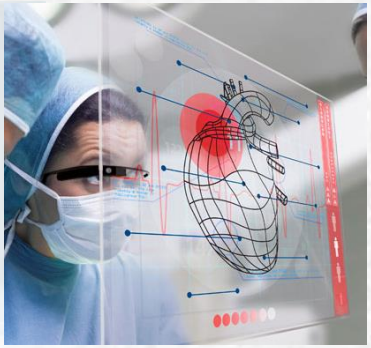

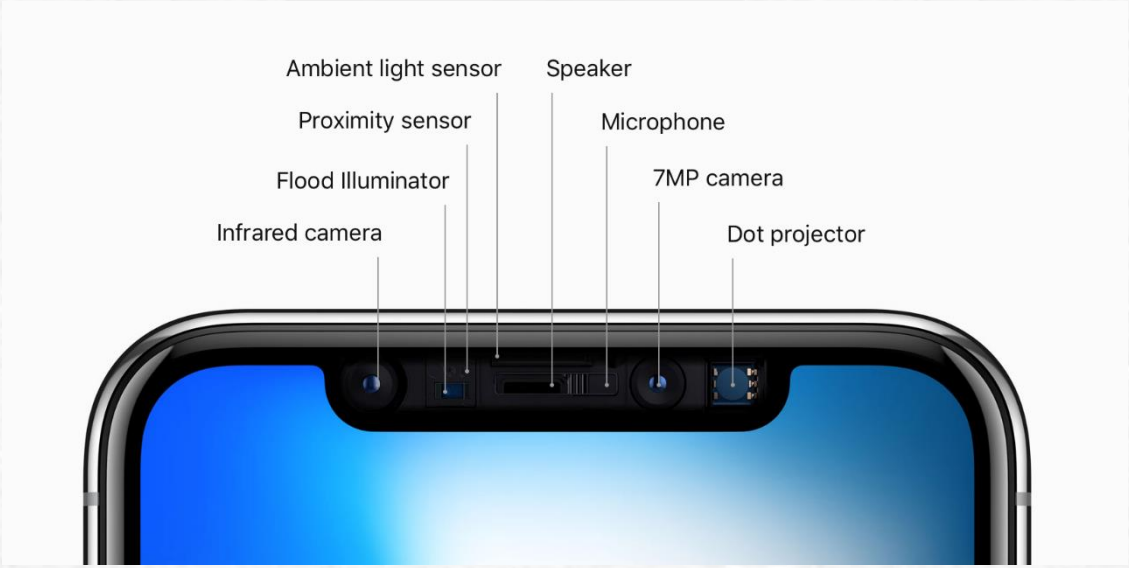

© WAYRAY (Reproduced with permission)

Applications – Augmenting reality

© Warner Bros. (Reproduced with permission)

Applications – Intelligent Devices

© Takeo Kanade, CMU

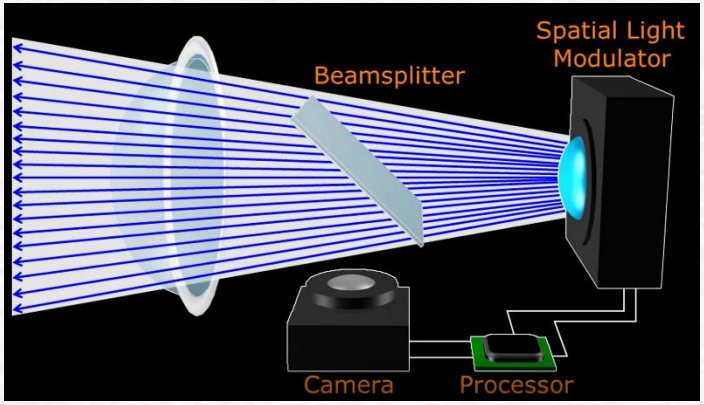

© Takeo Kanade, CMU

Applications – Self drive car

© Takeo Kanade, Robotics Institute, Carnegie Mellon University (Reproduced with permission)

The Ultimate Goal

"...apprehension by the senses supplies after all, directly or indirectly, the material of all human knowledge, or at least the stimulus necessary to develop every inborn faculty of the mind. It supplies the basis for the whole action of man upon the outer world... For there is little hope that he who does not begin at the beginning of knowledge will ever arrive at it's end“

[The Recent Progress of the Theory of Vision by H. v. Helmholtz (1873)]

"If our long-sought quest to create autonomous anthropomorphic automata is to succeed, we must first impart perceptual abilities to machines".

[A Guided Tour of Computer Vision, by V. Nalwa]

The goal of Computer Vision?

Ultimately emulating this…

So, how are we doing?

© NASA / JPL, reproduced with permssion

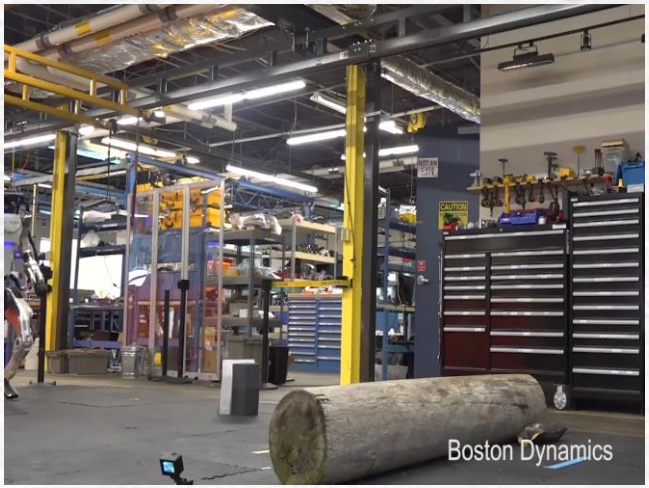

© Honda Motor Co. Ltd.
(Reproduced with permission)

Vision will…

- Solve harder problems with less constraints.

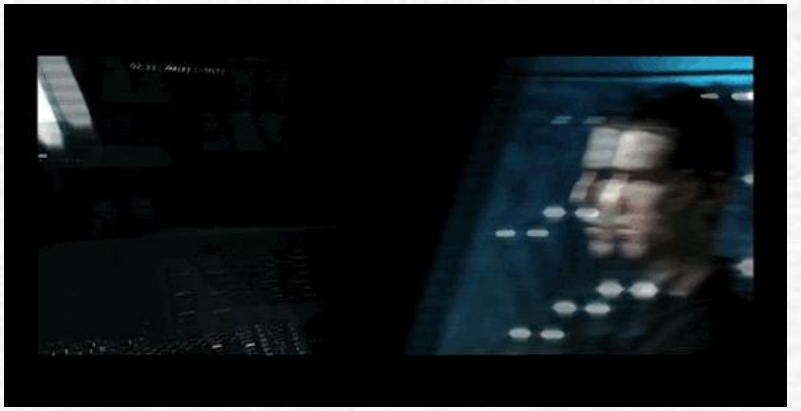

From http://www.telepresenceoptions.com/images/Minority_Report.gif

- Be a big part of User I/O
- Enhance security.
- Diagnose medical conditions reliably.
- Control vehicles automatically.
- Identify suspects through forensic analysis
- Provide more understanding about the environment.

Goals

- It is intended that the student should, at the end of the course...
  - Understand the broad subject of computer vision.
  - Understand the main algorithmic processes used to manipulate images.
  - Understand how information may be extracted from images, and the associated problems.
  - Be able to describe the various operations both algorithmically and mathematically.
  - Be able to code and test basic vision algorithms.
  - Be able to develop potential solutions to complex vision problems.

Based on *A Practical Introduction to Computer Vision with OpenCV* by Kenneth Dawson-Howe © Wiley & Sons Inc. 2014

Teaching Philosophy

- Flipped classroom
  - Prerecorded lectures to be studied in advance of class.
- Teaching by example.
- Hands-on experience of computer vision operations.
- Hands-on experience of coding computer vision operations
- Application of techniques worked on in group tutorials.
- Based on a single text +
  - RANSAC + PCA + SVM + K-means + Watershed +
  - Deep Learning + Stereo +…
- Lectures, slides & sample code provided online
  - On Blackboard (mymodule.tcd.ie)

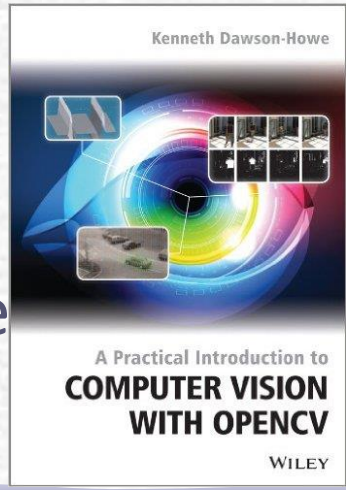

Based on *A Practical Introduction to Computer Vision with OpenCV* by Kenneth Dawson-Howe © Wiley & Sons Inc. 2014

Module delivery (most weeks)

- Around 4 pre-recorded lecture sessions (15-25 minutes each)
- 1 or 2 mini-tests to assess basic understanding of the pre-recorded material
  - Must be done before the Q&A session (or 0 marks).
- Q&A & 'advanced' topics session on Blackboard Collaborate to discuss any issues/questions relating to the pre-recorded lecture material
  - Students MUST have reviewed the pre-recorded lecture sessions. If you have not reviewed them then please do not ask any questions during the Q&A session.
  - Time permitting we will look at some more advanced/applied topics in Vision such as Ethics, Text recognition, Lane Detection, Recognition using Deep Learning, Face Recognition, Head pose estimation, etc.
- Tutorial session on Blackboard Collaborate to consider how to address real world problems using vision.
  - Students will be broken into small groups to come up with solutions (for 10-15 minutes), the course lecturer will visit the groups, and then when we come back together groups will present their solutions which will be discussed.
  - If the Q&A session does not have much to discuss, we may start the tutorial session during that hour.
- Discussion board

Module Assessment

- 10% Mini-tests which assess basic understanding of the pre-recorded material.
  - Must be done before the Q&A session (or 0 marks will be awarded).
  - Single attempt only.
- 30% Programming assignment.
  - Solving a real world vision problem and assessing how you have done.
- 10% Sample exam question(s) assignment.
  - Exam practice with feedback!
- 50% In person exam
  - 2 hours.
  - Some application of vision to real world problems type questions.
  - Some compare & contrast type questions.

Based on *A Practical Introduction to Computer Vision with OpenCV* by Kenneth Dawson-Howe © Wiley & Sons Inc. 2014

Module contents

Week	Beginning	Folder	Topic	Format	**Mini Tests (10%)**	Live Classes	Practical Component
		Introduction	What is Computer Vision	Live		Lecture Tuesday @2pm	
1	12/09/2022	Introduction	Module Organisation	Live		Lecture Wednesday @1pm	Install OpenCV & TIPS.
		Images	Camera Model & Digitisation	Pre-recorded			Experiment with TIPS
		Images	Colour Image models	Pre-recorded			**Assignment 1 (30%), Part 1**
2	19/09/2022	Images	Noise & Smoothing	Pre-recorded	Images before Tuesday @12noon	Q&A Session Tuesday @2pm	
		Binary	Thresholding	Pre-recorded		Tutorial Wednesday @1pm	
		Binary	Cleaning Binary Images	Pre-recorded	Binary before Tuesday @12noon		
		Regions	Connectivity & Connected Component Analy	Pre-recorded			**Assignment 1, Part 2**
3	26/09/2022	Regions	kMeans clustering	Pre-recorded	Regions before Tuesday @12noon	Q&A Session Tuesday @2pm	
		Geometric	Transformations	Pre-recorded		Tutorial Wednesday @1pm	
		Geometric	Interpolation	Pre-recorded	Geometric before Tuesday @12noon		
		Histograms	1D & 3D histograms	Pre-recorded			**Assignment 1, Part 3**
4	03/10/2022	Histograms	Equalisation, Comparison & Back Projection	Pre-recorded	Histograms before Tuesday @12noon	Q&A Session Tuesday @2pm	
		Video	Introduction	Pre-recorded		Tutorial Wednesday @1pm	
		Video	Background Models	Pre-recorded	Video before Tuesday @12noon		
		Edges	Edge detection	Pre-recorded			**Assignment 1, Part 4**
5	10/10/2022	Edges	Contour segmentation	Pre-recorded		Q&A Session Tuesday @2pm	
		Edges	Hough transform	Pre-recorded		Tutorial Wednesday @1pm	
		Edges	Least squared error & RANSAC	Pre-recorded	Edges before Tuesday @12noon		
		Learning & Evaluation	Learning & Ground Truth	Pre-recorded			**Assignment 1, Part 5**
6	17/10/2022	Learning & Evaluation	Performance	Pre-recorded	Learning & Evaluation before Tuesday @12noon	Q&A Session Tuesday @2pm	
		Recognition	SPR	Pre-recorded		Tutorial Wednesday @1pm	
		Recognition	SPR Features	Pre-recorded	Recognition before Tuesday @12noon		
		Regions II	Watershed segmentation	Pre-recorded			
		Regions II	Mean shift	Pre-recorded	Regions II before Tuesday @12noon	Q&A Session Tuesday @2pm	
8	31/10/2022	Recognition II	Template Matching	Pre-recorded		Tutorial Wednesday @1pm	
		Recognition II	Chamfer Matching	Pre-recorded			Submit Assignment 1
		Recognition II	SVM	Pre-recorded	Recognition II before Tuesday @12noon		- before 2-Nov @11:59pm
		Recognition III	Haar	Pre-recorded			**Assignment 2 (10%)**
9	07/11/2022	Recognition III	Haar	Pre-recorded		Q&A Session Tuesday @2pm	- Sample Exam Question
		Recognition III	PCA	Pre-recorded	Recognition III before Tuesday @12noon	Tutorial Wednesday @1pm	- released on 9-Nov @5:30pm
		Video II	Tracking - Mean Shift	Pre-recorded			
10	14/11/2022	Video II	Tracking - Mean Shift	Pre-recorded		Q&A Session Tuesday @2pm	Submit Assignment 3
		Video II	Tracking - Optical Flow	Pre-recorded	Video II before Tuesday @12noon	Tutorial Wednesday @1pm	- before 20-Nov @11:59pm
		Features	Corner detection - Moravec	Pre-recorded			
11	21/11/2022	Features	Corner detection - Harris	Pre-recorded		Q&A Session Tuesday @2pm	
		Features	SIFT Keypoints	Pre-recorded		Tutorial Wednesday @1pm	
		Features	SIFT Matching & Recognition	Pre-recorded	Features before Tuesday @12noon		
12	28/11/2022	Conclusions	How to address vision problems	Live		Lecture Tuesday @2pm	
		Conclusions	Sample exam paper tutorial	Live		Tutorial Wednesday @1pm	

Based on *A Practical Introduction to Computer Vision with OpenCV* by Kenneth Dawson-Howe © Wiley & Sons Inc. 2014